AMENDMENTS TO THE SPECIFICATION

Please amend the specification as follows:

On page 2, amend paragraph [0006] as follows:

[0006] Typically, and as shown in FIG. 1, DISAs 410, like local area networks (LANs) 420, and particularly LANs 420 connected to the Internet 430, transmit data using the Transmission Control Protocol/Internet Protocol (TCP/IP), see LAN connections 415 in FIG. 1. The TCP/IP protocol was designed for the sending of data across LAN-type architectures. However, DISAs 410, unlike LANs, contain a limited number of server nodes and are all generally located in very close proximity to one another. As such, DISAs 410 do not face much of the difficulties associated with transactions traveling over LANs 420, and as such, do not need much of the functionality and overhead inherent to the TCP/IP protocol. When DISAs are required to use TCP/IP, for example, and as shown by the solid line connections 415, such DISAs are disadvantaged by having to encapsulate and de-encapsulate data as it [[is]] travels within the cluster of servers. In fact, as the industry has provided LAN interconnects significantly larger than 100 Mb, i.e., 1 Gb and larger, both application and data resource servers have spent disproportionate amounts of Central Processing Unit (CPU) time processing TCP/IP communications overhead, and have experienced a negative impact in their price/performance ratio as a result. Therefore, although the use of TCP/IP protocol makes sense for transactions traveling across LANs, its use makes less sense for transactions traveling strictly within a DISA.

On page 2, amend paragraph [0007] as follows:

[0007] Briefly, an illustrative system provides an architecture and method of using a router node to connect a LAN to a server cluster arranged in a System Area Network (SAN). The router node is capable of distributing the LAN based traffic among the SAN server nodes. The LAN uses a LAN based protocol such as TCP/IP. The SAN uses a SAN based protocol such as Next Generation I/O (NGIO), Future I/O (FIO) or ~~INFINIBAND~~ INFINIBAND®. The illustrative system, unlike systems where SANs use a LAN based protocol, is able to achieve greater throughput by eliminating LAN based processing in portions of the system.

On page 3, amend paragraph [0009] as follows:

[0009] The cluster nodes contain a node management agent. The node management agent contains a session management agent and a policy management agent. These session management agents and policy management agents perform the cluster node portion of the same functionality as their ~~counter-parts~~ counterparts in the router node. One of the cluster nodes is selected as the management node and sets the policies on the router. The management node also includes an additional agent, the monitoring agent, which enables the management node to query the router node on a variety of statistics.

On page 4, amend paragraph [0011] as follows:

[0011]] As shown in FIGS. 2 and 3, the disclosed embodiments include all the functionality present in traditional DISA load balancing. However, unlike traditional DISAs that use the same protocols as the LANs they are connected to, i.e., TCP/IP, the disclosed embodiments instead use DISAs which operate under separate System Area Networks SAN based protocols. SAN based protocols are used in SAN-type architectures where cluster nodes are located in close proximity to one another. SAN based protocols provide high speed, low overhead, non-TCP/IP and highly reliable connections. By using such SAN based protocols DISAs are able to take advantage of the processing efficiencies associated with SAN based protocols such as NGIO, FIO and ~~INFINIBAND~~ INFINIBAND®, all of which are optimally suited for stand alone server clusters or SANs. This dual approach of having separate protocols for connected LANs and SANs allows the burden of the TCP/IP processing to be offloaded from application and data resource servers to router nodes which allows each type of node to concentrate on what it does best. Further, each of the different types of devices can be optimized to best handle the type of work they perform. The disclosed embodiments accommodate higher bandwidth TCP/IP processing than that found in traditional server networks.

On page 4, amend paragraph [0012] as follows:

[0012] As shown in FIGS. ~~2 and 4~~ 2-4, the Cluster or Server SAN Nodes 20, made up of application server nodes 220 and data resource server nodes 210, are connected to one another via a SAN 40. As shown in FIGS. 2-4, the SAN 40 in turn is connected to a Router Node 10. The Router Node 10 is thereafter connected to the LAN 30. Further, in greater detail as shown in

FIGS. 2-4, the Cluster Nodes 20 are attached to one or more Router Nodes 10 via a SAN 40. The Router Node 10 may be thereafter connected to a firewall 70 via a LAN 30, as shown in FIG. 3. Finally, the firewall 70 may be connected to the Internet 50 via a WAN 60 connection, as shown in FIG. 3. Other architectures connecting [[a]] SANs and LANs could also be used without departing from the spirit of the invention.

On pages 4-5, amend paragraph [0013] as follows:

[0013] FIG. 4 shows a detailed view of the disclosed embodiment. As shown, the Router Node 10 is connected at one end, to the LAN 30 through a LAN network interface controller (NIC) 170 using a TCP/IP connection, and at the other end, is connected through a SAN NIC 100 to the SAN 40 running a SAN based protocol such as NGIO, FIO or ~~INFINIBAND~~ INFINIBAND®. The Router Node 10 provides the translation function between the LAN protocol and the SAN protocol and distributes LAN originated communications across the Cluster Nodes 20. Also connected to the SAN 40 are Cluster Nodes 20. As a result, the SAN protocol is used for communication within the cluster and the LAN protocol is used for communication outside the cluster. Although the LAN and SAN protocols mentioned above can operate in conjunction with the disclosed embodiments, other LAN and SAN protocols may also be used without departing from the spirit of the invention.

On page 5, amend paragraph [0014] as follows:

[0014] Although only one Router Node 10 is depicted, it is contemplated that multiple Router Nodes 10 may be used (two Router Nodes 10 depicted in Fig. 4). If multiple Router Nodes 10 are used, they may be so arranged as to perform in a fail-over-type functionality, avoiding a single point of failure. In the fail-over-type functionality, only one Router Node 10 would be functioning at a time. But, if the node was to fail, the next sequential Router Node 10 would take over. Such an arrangement would provide protection against loosing communications for an extended period of time. Alternatively, if multiple Router Nodes 10 are used, they may be arranged such that they each work in parallel. If this parallel functionality were imposed, all of the Router Nodes 10 would be able to function at the same time. This architecture would likely allow greater throughput for the system as a whole since the data processing time to process TCP/IP packets that pass through a Router Node 10 is comparatively slow to the speed at which

the requests can be handled once reaching a SAN 40. Thus, in this architecture, enough Router Nodes 10 could be added to the system to balance the rate at which requests are received by the system (LAN activity) and the rate at which the system is able to process them (SAN activity).

On page 6, amend paragraph [0016] as follows:

[0016] The Cluster Nodes 20 include a Node Management Agent (NMA) 230. The NMA 230 further comprises a PMA 136, SMA 134 and a Monitoring Agent 236. Here, the PMA 136 and the SMA 134 perform similar functions to the corresponding agents in the Router Node 10, but do so for the Cluster Node 20. One or more of the Cluster Nodes 20 are designated as a Management Node 28 and sets policies on the Router Node 10. This Management Node 28 contains the only Cluster Node 20 with [[an]] a Monitoring Agent 236. The Monitoring Agent 236[[,]] provides the means to obtain various statistics from the Router Node 10. It may work with the PMA 136 to modify routing policy based on statistical information. Each Cluster Node 20 further includes a SAN Transport Protocol layer 202, SAN NIC Driver 201, and SAN NIC 200. The SAN NIC 200 is connected to the SAN 40.

On pages 6-7, amend paragraph [0018] as follows:

[0018] In the disclosed embodiment, the Router Node 10 is responsible for ensuring that the data from a Remote Client 80 connection gets consistently routed to the appropriate Cluster Node 20. The main purpose of Router Node 10, in acting as a bridge between the Remote Client 80 and a Cluster Node 20, is to handle the TCP/IP processing and protocol conversions between the Remote Client 80 and the Cluster Nodes 20. This separation of labor between Router Node 10 and Cluster Node 20 reduces processing overhead and the limitation otherwise associated with Ethernet rates. Further, the Router Node can be optimized in such a manner as to process its protocol conversions in the most efficient manner possible. In the same manner Cluster Nodes 20 can be optimized to perform its functions as efficiently as possible. In operation, the Router Node 10 probes the header field of incoming and outgoing packets to establish a unique connection between a remote client and a SAN Cluster Node 20. In the disclosed embodiment the set of Cluster Nodes 20 are viewed by Remote Clients 80 as a single IP address. This architecture allows the addition of one or more Cluster Nodes 20 in a manner that is transparent to the remote world. It is also contemplated that multiple IP addresses could be used to identify

the set of Cluster Nodes 20, [[and]] which would allow the reservation of a few addresses for dedicated virtual pipes with a negotiated quality of service.

On page 7-8, amend paragraph [0020] as follows:

[0020] In the Connection Establishment Phase, the Router Node 10 receives a request for connection from a Remote Client 80, and determines, based on connection information in the Policy Table, [[to]] which Cluster Node 20 to direct the request. FIG. 5 is an example of a Policy Table which comprises four fields: Service Type ("Services"), Eligibility, SAN Address and Weight. The Router Node 10 first determines, by probing the incoming TCP/IP packet, the type of service (service request type) for which the Remote Client 80 is requesting a connection. Based on the requested service, the Router Node 10 determines the type of authentication (authentication type) that is required for the requestor. The Eligibility field in the Policy Table encodes the type of authentication required for the service. The procedure to authenticate a requester may range from being a simple domain based verification to those based on encryption standards like Data Encryption Standard (DES), IP Security (IPSEC), or the like. Once the requester has been authenticated the eligible Cluster Nodes 20 capable of servicing the request are determined. Subsequently, one of these eligible Cluster Nodes 20 is selected based on the load balancing policy encoded for the particular service. The Weight field in the Policy Table contains a weighting factor that indicates the proportion of connection requests that can be directed to a particular Cluster Node 20 compared to other Cluster Nodes 20 for a given service. This Weight field is used by the load balancing routine to determine the Cluster Node 20 that would accept this request. Once the Cluster Node 20 has been identified to service the Remote Client 80, the Connection Establishment Phase is complete. The Router Node 10 then communicates with the Cluster Node 20 and completes the establishment of the connection.

On page 9, amend paragraph [0023] as follows:

[0023] The PMAs 136, existing on both the Router Nodes 10 and Cluster Nodes 20, and the RMA 130 and NMA 230 respectively, enable the Cluster Nodes 20 and Router Nodes 10 to inform and validate the services that each other expect to support. When the Cluster Node 20 is enabled, the PMA 136 on the Cluster Nodes' 20 Management Node 28 informs the Router Node 10, via entries in the Policy Table, see FIG. 3, [[of]] which services on what Cluster Nodes 20

are going to be supported. In addition, the Management Node 28 identifies the load-balancing policy that the Router Node 10 should implement for the various services. The load-balancing strategy may apply to all of the Cluster Nodes 20, or to a particular subset. The Management Node 28 is also involved in informing the Router Node 10 of any authentication policies associated with the services handled by the Cluster Nodes 20. Such authentication services (authentication types) may be based on service type, Cluster Node 20 or requesting Remote Client 80.